CUTTING A CAKE WITHOUT HARMING THE TOPPINGS

ARSENIY AKOPYAN[♠] AND EREL SEGAL-HALEVI[⋆]

Abstract. On a two-dimensional cake, there are m pairwise-disjoint toppings. It is required to cut the cake to $m+b$ pieces, such that each topping is entirely contained in a unique piece, and the total number of empty pieces (b, "blanks") is minimized. The minimal number of blanks naturally depends on the geometric constraints of the cake and its pieces. The paper present tight bounds on the number of blanks when the pieces must be connected, convex or rectangular.

arXiv:1604.00960v3 [cs.CG] 7 Nov 2016

Introduction

Let C ("cake" be a compact subset of the Euclidean plane $\mathbb{R}^2$. Let $Z_1, \ldots, Z_m$ be pairwise-disjoint subsets of C ("toppings"), such that $C \supsetneq Z_1 \cup \cdots Z_m$. It is required to create a partition of C, $C = Z'_1 \cup \cdots \cup Z'_{m+b}$, for some $b \geq 0$, such that:

$$\forall i \in 1 \ldots m : Z_i \subseteq Z'_i$$

i.e, each of the ms toppings is contained in a unique piece. Additionally, there may be b "blanks" — pieces with no topping.

A potential application is division of land. There is an initial division in which a subset of the land is divided among m people. We want to complete this division to a partition of the entire land, but we do not want to harm existing land-owners, so we allow only to extend the existing plots and (possibly) add blanks.

Subject to that requirement, it is desired to minimize the number of blanks. What is the smallest b that can be guaranteed, in the worst case, as a function of m?

Without further restrictions on the pieces, the answer is trivial: it is always possible to get $b = 0$ by setting e.g. $Z'_1 = Z_1 \cup C \setminus (Z_2 \cup \cdots \cup Z_m)$ (arbitrarily attach all the cake outside of the toppings to one of the toppings). The question becomes more interesting when there are geometric constraints on the pieces. A geometric constraint is represented by a family S of subsets of the plane. It is given that for each $i \in \{1, \ldots, m\}$, $Z_i \in S$. It is required that for each $i \in \{1, \ldots, m+b\}$, $Z'_i \in S$. Several such constraints are considered here:

(1) S is the family of path-connected sets, and the cake C is path-connected.
(2) S is the family of convex sets, and C is convex.
(3) S is the family of axes-parallel rectangles and C an axes-parallel rectangle.
(4) S is the family of axes-parallel rectangles and C a simply-connected rectilinear polygon.

A-priori, one could think that a sophisticated computational-geometric algorithm is needed to find a partition with a minimal b. However, it turns out that the following simple algorithm provides optimal worst-case guarantee:

> For each i, expand Z_i arbitrarily to the largest extent possible, while keeping it contained in C and an element of S and disjoint from the other toppings.

I.e, it does not matter how exactly the toppings are expanded, to what direction and in what order, as long as they are expanded to the maximum possible amount. Formally, say that

♠ Supported by People Programme (Marie Curie Actions) of the European Union's Seventh Framework Programme (FP7/2007-2013) under REA grant agreement n°[291734].

⋆ Supported by the Doctoral Fellowships of Excellence at Bar-Ilan University, the Mordecai and Monique Katz Graduate Fellowship Program at Bar-Ilan University.

a sequence $Z'_1, \ldots, Z'_m$, is *maximal*, if for each $i \in 1, \ldots, m$, there is no $Z^*_i \supsetneq Z'_i$ such that $Z^*_i \subseteq C$ and $Z^*_i \in S$ and $Z'_1, \ldots, Z'_{i-1}, Z^*_i, Z'_{i+1}, \ldots, Z'_m$ are pairwise-disjoint. The simple algorithm stated above can also be stated as:

(a) Find an arbitrary maximal set $Z'_1, \ldots, Z'_m$ such that for all i: $Z'_i \subseteq C$ and $Z'_i \in S$ and $Z'_i \supseteq Z_i$.
(b) Partition the reminder, $C \setminus (Z'_1 \cup \cdots \cup Z'_m)$, to a minimal number of elements of S; let that number be b and the parts $Z'_{m+1}, \ldots, Z'_{m+b}$.
(c) Return $Z'_1, \ldots, Z'_m, Z'_{m+1}, \ldots, Z'_{m+b}$.

This paper analyzes the performance of this algorithm for the four geometric scenarios presented above. During the analysis, for simplicity, it is assumed that the set $Z_1, \ldots, Z_m$ is already maximal. The analysis amounts to counting the blanks when the set of toppings is maximal.

The results are summarized in the table below. Some results are given in terms of the following function:

$$\text{DoubleRoot}(m) := \left\lfloor \sqrt{m-1} \right\rfloor + \left\lfloor \sqrt{m + \sqrt{m}} \right\rfloor.$$

Or, equivalently:

$$\text{DoubleRoot}(m) := \begin{array}{cc} 2k & \text{if } k^2 < m \leq k^2 + k \\ 2k+1 & \text{if } k^2 + k < m \leq (k+1)^2 \end{array}.$$

Thus,

$$\text{DoubleRoot}(m) = 2\sqrt{m} + O(1).$$

The expression in each row is the worst-case number of blanks b, as a function of the number of toppings m. In the case of a rectilinear polygon cake, the expression also depend on the number of reflex vertexes of the polygon, T.

Cake C	Pieces S	Blanks b in worst-case
Path-connected	Path-connected	0
Convex	Convex	$2m - 5$
Axes-parallel rectangle	Axes-parallel rectangles	$m - \text{DoubleRoot}(m)$
Simply-connected rectilinear polygon	Axes-parallel rectangles	$T + m - \text{DoubleRoot}(m)$

Note that each row actually represents two results: (a) An upper bound on the number of blanks, and (b) An example which shows that this upper bound is tight.

1. Path-Connected Cake and Path-Connected Pieces

In this section, S is the family of path-connected sets and the cake C is also path-connected. The following theorem is almost trivial; we prove it for completeness.

Theorem 1. *In a path-connected cake C with m path-connected toppings in a maximal arrangement, there are no blanks.*

Proof. Suppose by contradiction that there is a blank $H \subseteq C$. Since C is path-connected, H can be connected by a path to any of the toppings. In particular, there is some topping Z_i such that a path from H to Z_i through C does not pass through any other topping Z_j. The union of Z_i and the path and H is path-connected, contained in C and disjoint from the rest of the toppings, and it strictly contains Z_i — in contradiction to maximality. □

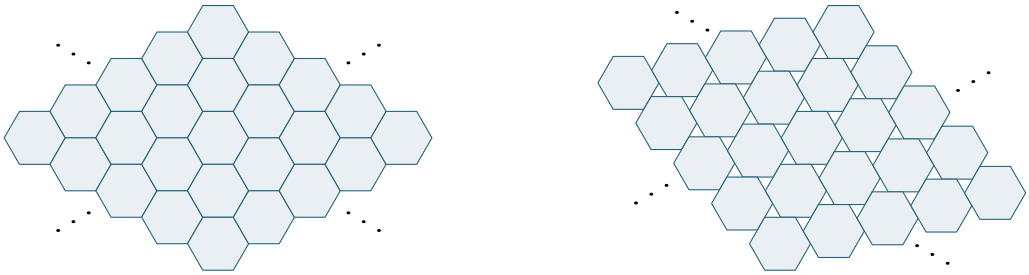

FIGURE 1. **Left:** a tiling of the plane by hexagons.
Right: a modification of the tiling where near each vertex there is a blank.

2. Convex Cake and Convex Pieces

In this section, S is the family of convex sets and the cake C is convex. It is proved that the worst-case number of blanks when the set of toppings is maximal is $2m-5$. Our techniques are very similar to [1].

2.1. **Lower bound.** To get intuition for the lower bound construction, consider the tiling of the plane with hexagons, shown in Figure 1/Left. Each point in which 3 hexagons meet is called a *vertex*. By slightly moving the hexagons, it is possible to create, near each vertex, a triangular blank, as is shown in Figure 1/Right. No hexagon can be expanded towards an adjacent blank while remaining convex and disjoint from the other hexagons, so the arrangement is maximal. Each blank touches three hexagons and each hexagon touches six blanks. Hence, the number of blanks is asymptotically twice the number of hexagons. In a finite tiling, there are boundary conditions. For example, when the tiling is contained in a square cake, there are $\Theta(\sqrt{m})$ triangles near the boundary of C. These can be discarded or joined with nearby toppings, so the total number of blanks is only $2m-\Theta(\sqrt{m})$.

Figure 2 shows a more sophisticated arrangement of toppings that has only a constant number of blanks in the boundary. The arrangement is based on a construction of Edelsbrunner, Robison and Shen [1]. For every integer $k \geq 0$, the toppings are:

- 3 large external hexagons;
- $3k$ medium intermediate hexagons;
- 3 or 4 or 5 or small inner polygons, depending on whether $m \mod 3$ is 0 or 1 or 2.

Near the vertexes of the toppings, there are triangular blanks. Each blank is triangular so it is adjacent to 3 toppings, so below each blank is counted three times:

- Each external hexagon is adjacent to 3 blanks, for a total of 9;
- Each intermediate hexagon is adjacent to 6 blanks, for a total of $18k$;
- The inner polygons are adjacent to:
  - When $m \mod 3 = 0$ — there are three inner quadrangles, which are adjacent to $3 \cdot 4 = 12$ blanks;
  - When $m \mod 3 = 1$ — there are one triangle and three pentagons, which are adjacent to $3 + 3 \cdot 5 = 18$ blanks;
  - When $m \mod 3 = 2$ — there are a triangle, a quadrangle, a pentagon and 2 hexagons, which are adjacent to $3 + 4 + 5 + 2 \cdot 6 = 24$ blanks.

All in all, the number of toppings and blanks is one of the following for every $k \geq 0$ (see Figure 2/Top):

- $m = 6 + 3k$ and $b = (9 + 18k + 12)/3 = 7 + 6k = 2m - 5$;
- $m = 7 + 3k$ and $b = (9 + 18k + 18)/3 = 9 + 6k = 2m - 5$;
- $m = 8 + 3k$ and $b = (9 + 18k + 24)/3 = 11 + 6k = 2m - 5$.

or one of the following (see Figure 2/Bottom):

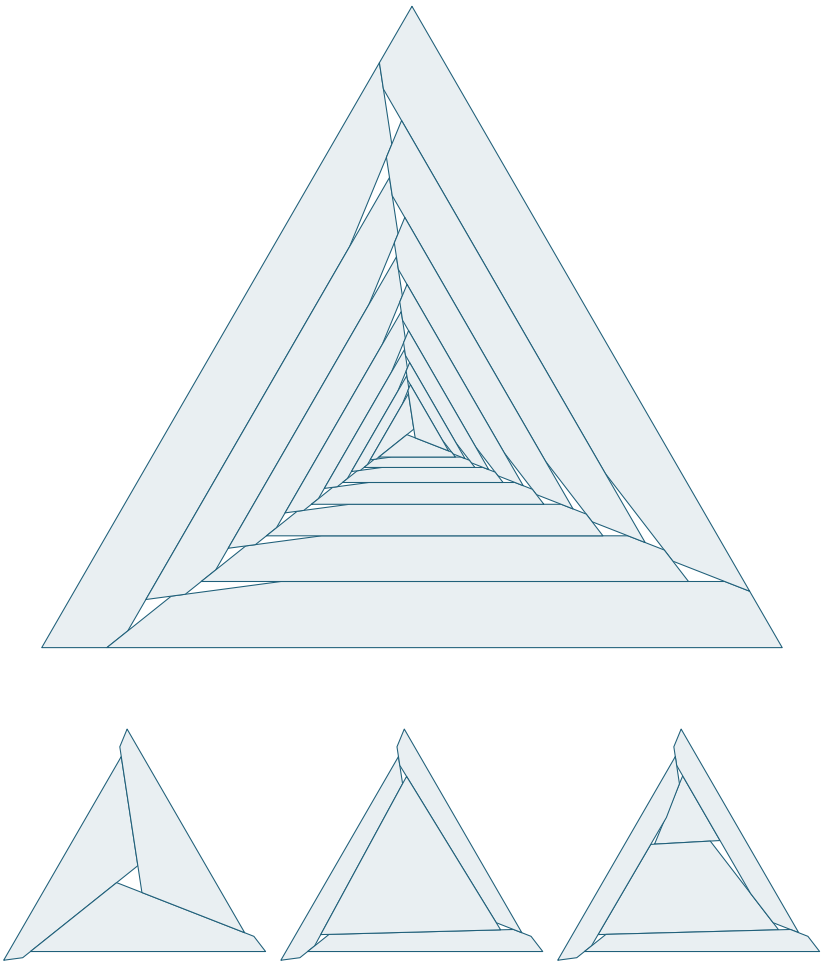

FIGURE 2. **Top:** Hexagonal toppings in a triangle cake.
Bottom: 3 or 4 or 5 toppings in the inner part of the cake, depending on $m \mod 3$.

- $m = 3$ and $b = 1 = 2m - 5$
- $m = 4$ and $b = 3 = 2m - 5$
- $m = 5$ and $b = 5 = 2m - 5$

In all cases the number of blanks is $2m - 5$, as claimed.

2.2. **Upper bound.** For the upper bound, we should distinguish between *holes* and *blanks*. A hole is a connected subset of C which is not contained in any of the toppings Z_i. A-priori, a hole need not be convex. We can imagine arrangements of convex toppings, such that the space between them is not convex. Hence, the number of blanks may be much larger than the number of holes (each convex component of a hole is counted as a different blank).

However, when the set of toppings is maximal, which is the case of interest here, it can be proved that all holes are convex. Hence, in particular, each hole is a single blank, and the number of blanks equals the number of holes. The proof is in Claim 2 of [4]:

Lemma 1. *When the set of toppings is maximal, every hole is a convex polygon.*

Theorem 2. *In a convex cake C with m convex toppings in a maximal arrangement, the number of blanks is at most $2m - 5$.*

Proof. Each side of a hole is formed by one of the Z_i-s. Let us call two toppings Z_i and Z_j *neighbors*, if they form two adjacent sides of some hole. Note that:

- To each hole correspond at least 3 neighbor-pairs — since each hole has at least 3 sides.

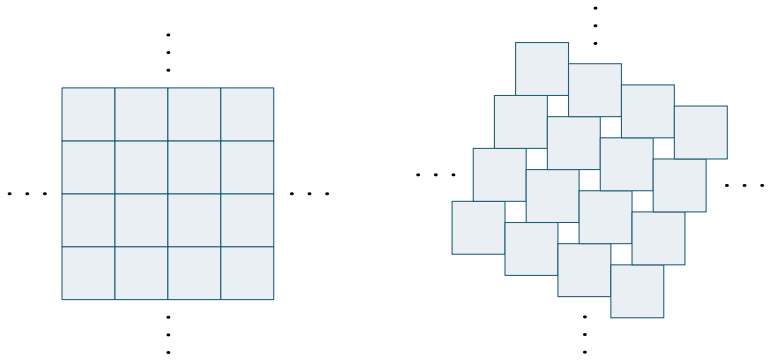

FIGURE 3. **Left:** a tiling of the plane by squares. **Right:** a modification of the tiling where near each vertex there is a blank.

- To each neighbor-pair Z_i, Z_j correspond at most two holes — since all such holes must have a side co-linear with the segment in which the boundaries of Z_i and Z_j intersect.
- Moreover, there are at least 3 neighbor-pairs near the boundary of C. This is because there are at least 3 toppings adjacent to the boundary of C (since if the boundary of C belongs only to one or two toppings, these topping/s must fill the entire cake and leave no room for holes). These neighbor-pairs correspond to at most one hole, since there is no hole in the side adjacent to the boundary.

Therefore:

$$\#\text{holes} \leq (\#\text{pairs} - 3) \cdot 2/3 + 3 \cdot 1/3 = \#\text{pairs} \cdot 2/3 - 1.$$

The "neighbor" relation defines a planar graph which is denoted by G. Denote by V the set of vertexes of G — which are the toppings $Z_1, \ldots, Z_m$. Denote by E the set of edges of G — which are the neighbor-pairs defined above. Denote by F the set of faces of G. By Euler's formula: $|V| - |E| + |F| = 2$. Each face is at least a triangle so it is adjacent to at least three edges. Each edge is adjacent to at most two faces. Hence, $|F| \leq 2|E|/3$. Substituting into Euler's formula gives $|V| - |E|/3 \geq 2$ so $|E| \leq 3|V| - 6$. Hence, the number of neighbor-pairs is at most $3m - 6$. Hence, the number of holes is at most $2m - 5$, and so is the number of blanks: $b \leq 2m - 5$. □

3. Rectangular Cake and Rectangular Pieces

In this section, S is the family of axes-parallel rectangles and the cake C is also an axes-parallel rectangle. It is proved that the worst-case number of blanks when the set of toppings is maximal is $m - 2\sqrt{m} - O(1)$.

3.1. **Lower bound.** Consider the tiling in Figure 3/Left. Here, four squares meet near each vertex. By moving them slightly, we get the arrangement in Figure 3/Right, where near each vertex there is a square blank. Each blank touches four squares and each square touches four blanks. Hence, the number of squares and blanks is asymptotically the same.

When the tiling is finite, there are boundary conditions. Suppose the cake is a square and it is tiled by $m = (k+1)^2$ smaller squares in a grid of $k+1$ times $k+1$. Then, the number of inner vertexes, that can be converted to blanks, is $k^2 = m - 2k - 1 = m - \text{DoubleRoot}(m)$.

When m is not a square number, e.g. $m = (k+1)^2 + t$, the lower bound can be constructed by gluing $t-1$ squares and one rectangle at the right and then at the top of the $(k+1) \times (k+1)$ tiling. Every additional rectangle adds a 4-vertex that can be converted to a blank, except the first rectangle at the right and the first rectangle at the top. Therefore the number of blanks is again $m - \text{DoubleRoot}(m)$.

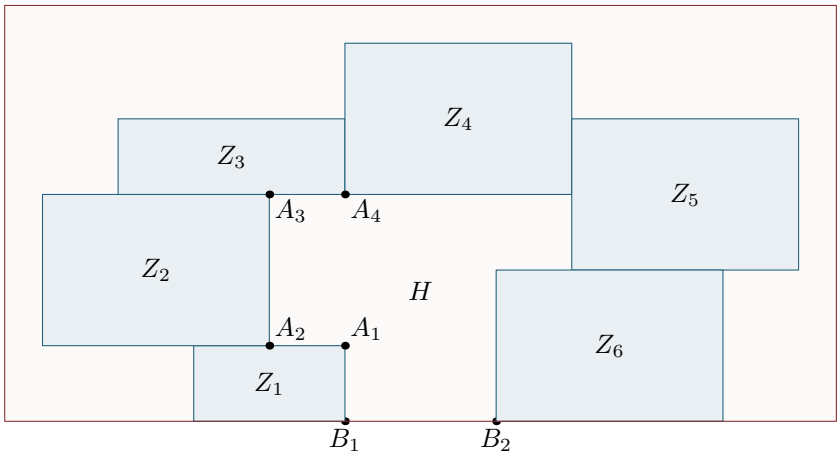

FIGURE 4. A rectangular cake C, rectangular toppings Z_j, and a hole H. Note the configuration is not maximal.

3.2. **Upper bound part A: All holes are axes-parallel rectangles.** Similarly to the convex setting, we prove that every hole is an axes-parallel rectangle, and thus every hole corresponds to a single blank.

Let H be a hole (a part of C not contained in any of the toppings Z_i). The boundary of H is made of line-segments that are parts of the boundaries of the Z_i, and possibly some line-segments that are parts of the boundary of C. The examples below refer to Figure 4.

Definition 1. *A* vertex of H *is a point on H's boundary that is a vertex of Z_i or of C.*

Note that this includes the vertexes of H as a polygon, but may also include "flat vertexes" (vertexes with angle 180°, such as the vertex A_4).

Definition 2. *An* edge of H *is a line-segment between two adjacent vertexes of H.*

Definition 3. *A vertex/edge that is contained in the boundary of C is called a* boundary-vertex/edge. *Otherwise it is called an* inner-vertex/edge.

Definition 4. *A* connection-vertex *is a boundary-vertex connecting an inner-edge and a boundary-edge (e.g. the vertexes B_1, B_2).*

By our definition of the vertexes of a hole H, it is clear that each edge of H is contained either in an edge of a single topping Z_i (an inner edge) or in an edge of the cake C (a boundary edge).

Definition 5. *For each edge e, let $Con(e)$ be the (unique) rectangle one of whose edges contains e. This may be either a single Z_i or C.*

Definition 6. *Given a vertex v and an edge e of a hole H, say that v* is exposed in *e if v is a vertex of $Con(e)$.*

For example, the vertex A_3 is exposed in the edge A_2A_3 but not in the edge A_4A_3.

Definition 7. *A vertex of a hole H is called:*

- convex — *if the internal angle adjacent to it (the angle through the interior of H) is* 90°.
- flat — *if the internal angle adjacent to it is* 180° *(e.g. vertex A_4).*
- reflex — *if the internal angle adjacent to it is* 270° *(e.g. vertex A_1).*

The following lemmas all assume that the configuration of toppings is maximal.

Lemma 2. *Every inner edge of a hole cannot have both of its vertexes exposed.*

Proof. Let e be an inner edge of a hole H. Then $Con(e)$ must be one of the toppings Z_i. Since both edges of e are exposed, e is equal to an edge of Z_i. But then, the rectangle Z_i can be stretched over e and into the hole H, contradicting the maximality condition. □

For example, the inner edge A_2A_3 has both its vertexes exposed, and indeed the rectangle Z_1 can be stretched rightwards into the hole and the configuration is not maximal.

Lemma 3. *Every inner vertex is exposed in at least one of its two adjacent edges.*

Proof. Let v be an inner vertex and let e, e' be its two adjacent edges (which must both be inner edges). Suppose by contradiction that v is not exposed in e nor in e'. Then v is both internal to an edge of $Con(e) = Z_i$ and internal to an edge of $Con(e') = Z_j$. But this implies that Z_i and Z_j overlap, contradicting the pairwise-disjointness condition. □

Lemma 4. *If an inner vertex is either reflex or flat (Definition 7), then it is exposed in both its edges.*

Proof. An inner reflex-vertex of a hole is necessarily a vertex of some topping Z_i which does not touch any other topping. Hence, its two adjacent edges are edges of Z_i, so it is exposed in both of them (see e.g. vertex A_1).

An inner flat-vertex of a hole is necessarily a vertex of two adjacent toppings (see e.g. vertex A_4) so it is exposed in both of them. □

Lemma 5. *Every connection-vertex (Definition 4) is exposed in its adjacent inner edge.*

Proof. Since the cake C is a rectangle, a connection-vertex of a hole can happen only when the vertex of a rectangle-topping touches the interior of an edge of C. □

Lemma 6. *Every vertex of a hole must be inner.*

Proof. By contradiction, let H be a hole with a boundary-vertex. Let $e_1, \ldots, e_t$ be a maximal path of edges of H that are inner-edges. Let v_0 be the vertex before e_1 and v_t the vertex after e_t; note that v_0 and v_t are connection-vertexes (e.g. B_1, B_2).

By Lemma 5, v_0 is exposed in e_1 and v_t is exposed in e_t. So, in the directed path from v_0 to v_t through $e_1, \ldots, e_t$, the first vertex is exposed in the edge after it (e.g. B_1 is exposed in B_1A_1) and the last vertex is exposed in the edge before it (e.g. B_2). All other vertexes in the path are inner-vertexes, so by Lemma 3 each of them is exposed in at least one edge. All in all, there are $t+1$ exposed vertexes but only t edges, so by the pigeonhole principle, at least one edge in the path must have two of its vertexes exposed — a contradiction to Lemma 2. □

Lemma 7. *No vertex of a hole is reflex or flat.*

Proof. By contradiction, let v be a reflex or flat vertex of a hole H (e.g. the reflex vertex A_1). By Lemma 6 it is an inner-vertex so by Lemma 4 it is exposed in both its adjacent edges.

Consider the cyclic path starting from v and going over the edges and vertexes of H until it reaches v again (e.g. from A_1 through A_1A_2 to A_2 through A_2A_3 to A_3 ... to B_1 through B_1A_1 to A_1). In this path, v is exposed both in the edge after it (e.g. A_1A_2) and in the edge before it (e.g. B_1A_1). By Lemma 6, all vertexes in this path are inner-vertexes, so by Lemma 3, each of them is exposed in at least one edge. All in all, there are $t+1$ "exposed" relations but only t edges, so by the pigeonhole principle, at least one edge in the path must have two of its vertexes exposed — a contradiction to Lemma 2. □

The latter two lemmas yield:

Corollary 1. *Every hole H must be a rectangle with no vertices on its edges, and must be strictly in the interior of C (not adjacent to its boundary).*

So all holes in a maximal arrangement look like Figure 3: each hole is a rectangle adjacent to four toppings.

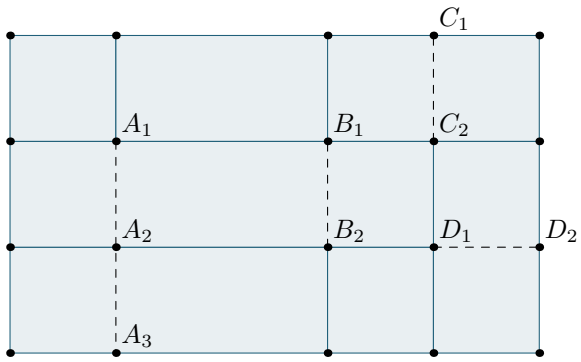

FIGURE 5. Bounding the number of 4-vertices in a rectangle tiling.

3.3. **Upper bound part B: Counting the holes.** Our upper bound matches the lower bound of Subsection 3.1: $m - \text{DoubleRoot}(m)$, where:

$$\text{DoubleRoot}(m) := \begin{matrix} 2k & \text{if } k < m \leq k^2 + k \\ 2k+1 & \text{if } k^2 + k < m \leq (k+1)^2 \end{matrix}.$$

The bound relies on the following lemma:

Lemma 8. *For every non-negative integers k_1, k_2, t:*

$$(k_1 - 1)(k_2 - 1) - t \leq k_1 k_2 - t - \textit{DoubleRoot}(k_1 k_2 - t).$$

Proof. It is enough to show that $\text{DoubleRoot}(k_1 k_2) \geq k_1 + k_2 - 1$. For fixed sum $k_1 + k_2$ the product maximize when summands are equal or have difference 1. In both cased we have the equality. □

Suppose a rectangle is partitioned to m smaller rectangles (with no holes). Define a 3-*vertex* as a point in which three rectangles meet and a 4-*vertex* as a point in which four rectangles meet.

Lemma 9. *Suppose a rectangular cake C is partitioned to m rectangles. Then, the number of 4-vertices is at most $m - \textit{DoubleRoot}(m)$.*

Proof. First, suppose that there are only 4-vertices (no 3-vertices). Then C is partitioned, by the lines containing all sides of all rectangles, to a grid of k_1 by k_2 smaller rectangles. So the total number of pieces is $k_1 k_2$ and the total number of 4-vertices is $(k_1 - 1)(k_2 - 1)$. Then the lemma follows from Lemma 8.

Next, suppose that there are 3-vertices (refer to Figure 5 for the examples). Choose one 3-vertex (e.g. A_1). Add a segment that separates the rectangle that does not have vertex at v to two parts (e.g. the segment A_1A_2). This converts the 3-vertex to 4-vertex. Additionally, there are several cases depending on where the other end of the additional segment lands:

(a) If it lands on the boundary of C (like with D_1D_2 and C_2C_1), then no further action is required; the number of rectangles grows by 1 and the number of 4-vertices grows by 1.

(b) If it lands on the boundary of another rectangle (like with A_1A_2), then a new 3-vertex is created, and can be handled in the same way by continuing the segment. This must stop because eventually the line hits the boundary of C.

(c) If it lands on a 4-vertex (like with B_1B_2), then an additional 4-vertex is created; the number of rectangles grows by 1 and the number of 4-vertices grows by 2.

In all cases, the number of 4-vertices grows at least as much as the number of rectangles.

Continue this procedure until all 3-vertices disappear. Then, the partition is a grid formed by the lines containing the sides of the all original rectangles. Suppose the grid has $k_1 \times k_2$ rectangles (e.g. in Figure 5, $k_1 = 3$ and $k_2 = 4$). So, the total number of rectangles is $k_1 k_2$

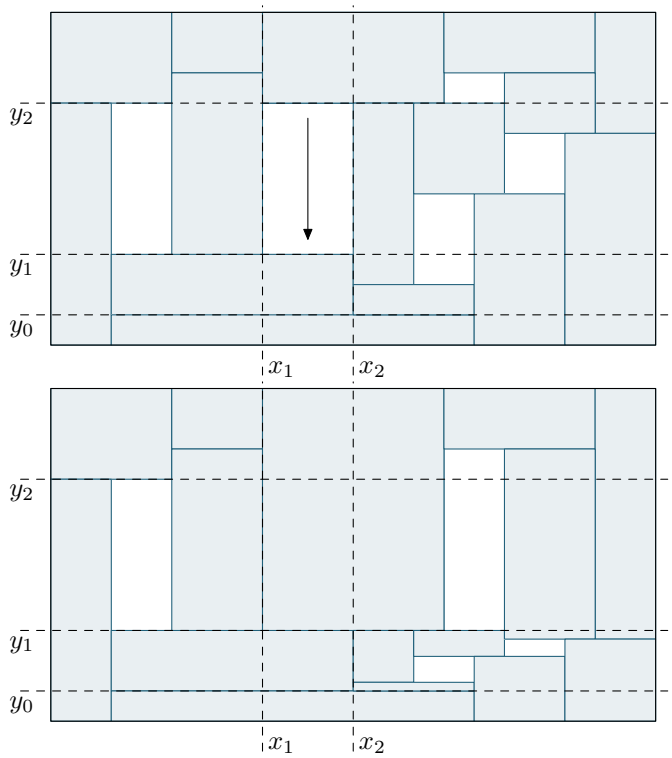

FIGURE 6. **Top:** An arrangement with 4 holes. **Bottom:** One of the holes is contracted vertically. The other rectangles and holes change their size but remain in the picture.

and the total number of 4-vertices is $(k_1 - 1)(k_2 - 1)$. Suppose we had t removal steps. So, the original number of rectangles was $m = k_1 k_2 - t$ and the original number of 4-vertexes was at most $(k_1 - 1)(k_2 - 1) - t$. From Lemma 8 it follows that the number of original 4-vertices is at most $m - \text{DoubleRoot}(m)$. □

We now use Lemma 9 to our general problem, regarding the number of holes.

Theorem 3. *In a rectangular cake C with m rectangular toppings in a maximal arrangement, the total number of holes is at most $m - \text{DoubleRoot}(m)$.*

Proof. We show that the holes can be removed one by one, by contracting each hole to a 4-vertex without changing the number of rectangles.

Each hole is contracted in two steps: *vertical* and *horizontal*.

The vertical contraction is illustrated in Figure 6. Consider the hole $[x_1, x_2] \times [y_1, y_2]$. As explained in the previous subsection, there are two options regarding the rectangles surrounding the hole: either the top rectangle has a side at $x = x_1$ and the bottom rectangle has a side at $x = x_2$, or vice versa. W.l.o.g, we assume that the former is the case, as in Figure 6/Top.

Let y_0 be the bottom coordinate of the rectangle below the hole. Transform the arrangement of rectangles in the following way. Every point (x, y), where $x > x_2$ and $y \in [y_0, y_2]$, is transformed to (x, y'), where:

$$y' := y_0 + \frac{y_1 - y_0}{y_2 - y_0} \cdot (y - y_0),$$

so the ray $(x > x_2, y = y_2)$ goes to $(x > x_2, y = y_1)$ and the ray $(x > x_2, y = y_0)$ remains in its place. Now, the rectangle on top of the hole can be extended down, so that it covers the hole (see 6/Bottom).

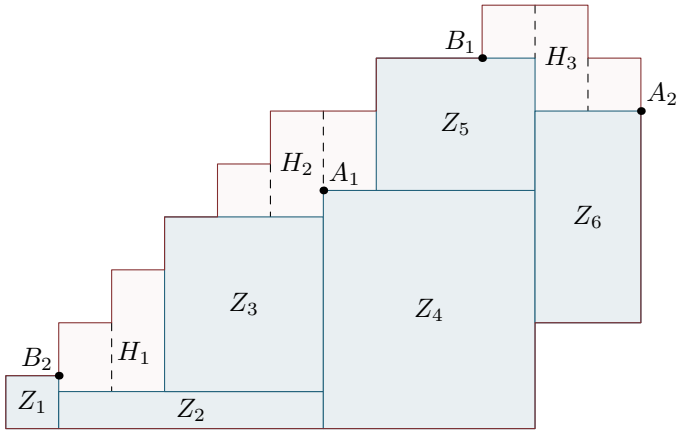

FIGURE 7. A rectilinear cake with six rectangular topping $Z_1, \ldots, Z_6$, and three rectilinear holes $H_1, \ldots, H_3$. Note the configuration is not maximal since Z_6 can be extended upwards into H_3. The dashed lines inside the holes indicate a possible partitioning of the holes to rectangular blanks.

The horizontal contraction is very similar. Let x_0 be the left coordinate of the rectangle to the left of the hole. Every point (x, y), where $y < y_1$ and $x \in [x_0, x_2]$, is transformed to (x', y), where:

$$x' := x_0 + \frac{x_1 - x_0}{x_2 - x_0} \cdot (x - x_0),$$

so the ray $(y < y_1, x = x_2)$ goes to $(y < y_1, x = x_1)$ and the ray $(y < y_1, x = x_0)$ remains in its place. Now, instead of the hole, there is a single 4-vertex at the point (x_1, y_1).

The shrinking does not change the combinatorics of the arrangement: the number of toppings does not change, and no new holes are created.

After all holes are contracted, the situation is as in Lemma 9, where the number of 4-vertices is upper-bounded. Therefore, the number of holes in the original configuration is upper-bounded by the same expression. □

4. Rectilinear-polygonal cake and rectangular pieces

In this section, S is still the family of axes-parallel rectangles, but now the cake C can be a simply-connected axes-parallel rectilinear polygon. The "complexity" of a rectilinear polygon is characterized by the number of its *reflex vertexes* — vertexes with internal angle 270°. It is known that a rectilinear polygon with T reflex vertexes can always be partitioned to at most $T + 1$ rectangles [3, 2], and this bound is tight when the vertexes of C are in general position. Since our goal is to bound the number of blank rectangles — it is expected that the bound should depend on T, in addition to m (the number of toppings).

4.1. **Lower bound.** Consider the tiling of Figure 3 in Subsection 3.1. Assume that the tiling is contained in a large square, and the square is connected in a narrow pipe to a rectilinear polygon with T reflex vertexes. As explained there, the number of blanks in this tiling is $m - \text{DoubleRoot}(m)$. The connected rectilinear polygon is a single *hole*, but it may contain as many as $T + 1$ *blanks* (its rectangular components). All in all, there are $m + T - \text{DoubleRoot}(m) = m + T - 2\sqrt{m} - O(1)$ blanks.

4.2. **Upper bound part A: characterizing the holes.** Recall that we still assume that the configuration of toppings is maximal. The first three lemmas of Subsection 3.2 (Lemmas 2 and 3 and 4) are still true since they do not depend on the shape of C. But the last three lemmas are not true.

The difference starts at Lemma 5: since C is not a rectangle, there may be connection-vertexes which are vertexes of C and not of any topping Z_i; see Figure 7. We distinguish between these cases as follows:

Definition 8. *(a) A* topping-connection-vertex *is a connection-vertex that is a vertex of a topping (e.g. A_2, B_2);*

(b) A cake-connection-vertex *is a connection-vertex that is a vertex of the cake (e.g. B_1, B_2).*

Lemma 5 is now replaced by:

Lemma 10.

(a) A topping-connection-vertex is exposed in its adjacent inner edge;

(b) A cake-connection-vertex is a reflex vertex of C.

Proof. Let v be a connection-vertex. By Definition 4, v is on the boundary of C, and it connects an inner edge with a boundary edge. Now:

(a) If v is a vertex of a topping, then it is exposed in the edge of that topping, which is an inner edge; examples are $A_@$ and B_2.

(b) If v is a vertex of the cake, then it must be a reflex vertex of the cake. This is because a non-reflex vertex necessarily connects two edges of the cake, so it cannot be a connection-vertex. Examples are B_1 and B_2. □

Lemma 6 is also no longer true; a hole may be adjacent to the boundary of C. That lemma is replaced by two lemmas:

Lemma 11. *Each hole is simply-connected.*

Proof. Suppose by contradiction that a hole is not simply-connected. So there is a "gap" inside the hole. Since the cake itself is simply-connected, this "gap" must be made of toppings. But these toppings can be extended towards the hole — contradicting the maximality assumption. □

Hence, each hole has a connected boundary, so we can talk about "*the* boundary" of a hole.

Lemma 12. *If a hole is adjacent to the boundary of C, then it must have at least one cake-connection-vertex.*

Proof. If the boundary of a hole is adjacent to the boundary of C, then it must have at least two connection-vertexes. Choose two neighborhood connecting-vertexes, connected by a path of inner edges.

If both of them are topping-connection-vertexes, then by Lemma 10, both must be exposed in the inner edges adjacent to them. But then, similarly to the proof of Lemma 6, the path between these vertexes must contain an edge both of whose vertexes are exposed — a contradiction to Lemma 2.

Therefore, at least one connection-vertex must not be a topping-connection-vertex, so it must be a cake-connection-vertex. □

Lemma 7 is also no longer true. A hole may have reflex vertexes. E.g, in H_2 there are two reflex vertexes: one of them (A_1) is a reflex-inner-vertex, and the other is a reflex-boundary-vertex (it is a reflex vertex of C). Similarly, a hole may have flat vertexes. E.g, in H_3, the vertex B_2 is flat. So lemma 7 is replaced by:

Lemma 13. *Let v be reflex/flat-inner-vertex of a hole H. Then, any path on the boundary of H starting at v, either clockwise or counter-clockwise, must meet zero or more convex-inner-vertexes, and then a cake-connection-vertex.*

Proof. Lemma 4 is still true: v is exposed in both its adjacent edges. Hence, regardless of the direction of the path, v is exposed in the edge following it in the path.

We now enumerate the possible cases:

(a) Suppose that, after meeting zero or more convex-inner-vertexes, we meet a reflex/flat-inner-vertex of H (possibly the vertex v itself). Then, since that vertex is exposed in the edge before it in the path, the path must contain an edge both of whose vertexes are exposed — a contradiction to Lemma 2.

(b) Suppose that, after meeting zero or more convex-inner-vertexes, we meet a topping-connection-vertex (like A_2 or B_2). Then this vertex is exposed in the edge before it in the path, which again leads to a contradiction to Lemma 2.

(c) The only remaining option is that, after meeting zero or more convex-inner-vertexes, we meet a cake-connection-vertex. This is the only kind of vertex which may be not-exposed. □

As an example, suppose we are on vertex A_1, and start to travel either clockwise or counter-clockwise; in either case we meet a reflex-vertex of C, which is not exposed in the edge between it and A_1.

4.3. **Upper bound part B: Counting the holes.** Every internal hole is a rectangle. This follows directly from the proofs in Section 3, since when considering internal holes, we can imagine that the actual cake is C' — a large rectangle that bounds our rectilinear-polygon C — and that $C' \setminus C$ is filled with toppings. This has no effect on the internal holes. Therefore, the upper bound in Subsection 3.3, namely $m - \text{DoubleRoot}(m)$, is still a valid upper bound on the number of internal rectangular blanks.

To bound the total number of blanks, we have to handle the blanks resulting from boundary holes. The main step is to bound the number of reflex vertexes in the boundary holes.

Lemma 14. *Let B_h be the set of boundary holes, T_h the set of reflex vertexes in all boundary holes, and T_C the set of reflex vertexes of C. Then:*

$$|B_h| + |T_h| \leq |T_C|.$$

Proof. The lemma is proved by the following 1-1 mapping from $B_h \cup T_h$ into T_C.

- Each element of B_h (a boundary hole) is mapped to a cake-connection-vertex, which exists by Lemma 12. This vertex is in T_C, by Lemma 10(b).
- Each element $v \in T_h$ which is an inner-vertex, corresponds to two cake-connection-vertexes — one at the clockwise path and one at the counter-clockwise path. At most one of them is already mapped to the hole containing v, so v can be mapped to the other cake-connection-vertexes.

  Note that, by Lemma 13, the paths from v cannot meet any other inner-reflex-vertex. Hence, there is at most one inner-reflex-vertex per hole. So different v are mapped to different cake-connection-vertices.
- Each element $v \in T_h$ which is a boundary-vertex, is a reflex-vertex of C, so it is mapped to itself. Note that v is not a connection-vertex, since its two adjacent edges are boundary-edges (edges of C). So it was not mapped in the previous two bullets.

□

Theorem 4. *In a simply-connected rectilinear cake with T reflex vertexes and m rectangular toppings in a maximal arrangement, the total number of holes is at most $T + m - \text{DoubleRoot}(m)$.*

Proof. By Lemma 14, the total number of reflex-vertexes in all boundary holes satisfies:

$$|T_h| \leq T - |B_h|,$$

where $T = |T_C|$ = the number of reflex vertexes of C.

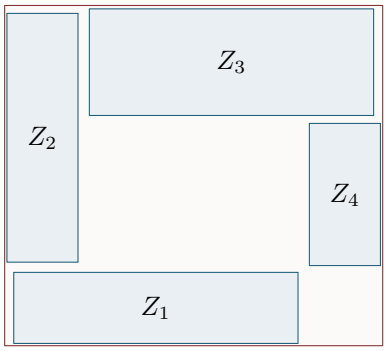

FIGURE 8. A configuration of toppings that can be extended in two ways: optimal (Z_4 leftwards) and sub-optimal (Z_4 downwards).

A rectilinear polygon with T reflex vertexes can always be partitioned to at most $T+1$ rectangles [3, 2]. Hence, the boundary holes in the set B_h can be partitioned to rectangles such that the total number of rectangles is at most $(T-|B_h|)+|B_h| = T$.

Adding the $m-\text{DoubleRoot}(m)$ internal blanks gives that the total number of (rectangular) blanks is at most: $T+m-\text{DoubleRoot}(m)$. □

5. Conclusion

We proved that, if a given configuration of toppings is extended to an arbitrary maximal configuration, then the number of blanks equals the worst-case minimum given the number of toppings.

However, in some cases it may possible to extend the toppings in a way that attains a smaller number of blanks — less than the worst-case minimum. For example, in Figure 8, if Z_4 is extended downwards, then we get one blank, which is the worst-case minimum for $m=4$ toppings. However, if Z_4 is extended leftwards, then we can get zero blanks.

The problem of finding the smallest number of blanks in each specific configuration is left for future work.

6. Acknowledgments

This paper started as a question in the Math Overflow website. We are grateful to Reuven Cohen, David Peleg, Herbert Edelsbrunner and Rom Pinchasi for discussions about the problem studied in this paper.

References

[1] H. Edelsbrunner, A. D. Robison, and X.-J. Shen. Covering convex sets with non-overlapping polygons. *Discrete Mathematics*, 81(2):153–164, Apr. 1990.

[2] D. Eppstein. Graph-Theoretic Solutions to Computational Geometry Problems. In C. Paul and M. Habib, editors, *Graph-Theoretic Concepts in Computer Science*, volume 5911 of *Lecture Notes in Computer Science*, pages 1–16. Springer Berlin Heidelberg, 2010.

[3] J. M. Keil. Polygon Decomposition. In *Handbook of Computational Geometry*, pages 491–518. University of Saskatchewan Saskatoon, Sask., Canada, 2000.

[4] R. Pinchasi. On the perimeter of k pairwise disjoint convex bodies contained in a convex set in the plane. *Combinatorica*, pages 1–27, July 2015.

Arseniy Akopyan, Institute of Science and Technology Austria (IST Austria), Am Campus 1, 3400 Klosterneuburg, Austria
E-mail address: akopjan@gmail.com

Erel Segal-Halevi, Bar-Ilan University, Ramat-Gan 5290002, Israel
E-mail address: erelsgl@gmail.com